arXiv:1510.05502v2 [cs.DM] 13 Jul 2016

The complexity of signed graph and edge-coloured graph homomorphisms

Richard C. Brewster*† Florent Foucaud‡ Pavol Hell§†¶ Reza Naserasr‖

July 14, 2016

Abstract

We study homomorphism problems of signed graphs from a computational point of view. A signed graph (G, Σ) is a graph G where each edge is given a sign, positive or negative; $\Sigma \subseteq E(G)$ denotes the set of negative edges. Thus, (G, Σ) is a 2-edge-coloured graph with the property that the edge-colours, $\{+, -\}$, form a group under multiplication. Central to the study of signed graphs is the operation of switching at a vertex, that results in changing the sign of each incident edge. We study two types of homomorphisms of a signed graph (G, Σ) to a signed graph (H, Π): ec-homomorphisms and s-homomorphisms. Each is a standard graph homomorphism of G to H with some additional constraint. In the former, edge-signs are preserved. In the latter, edge-signs are preserved after the switching operation has been applied to a subset of vertices of G.

We prove a dichotomy theorem for s-homomorphism problems for a large class of (fixed) target signed graphs (H, Π). Specifically, as long as (H, Π) does not contain a negative (respectively a positive) loop, the problem is polynomial-time solvable if the core of (H, Π) has at most two edges, and is NP-complete otherwise. (Note that this covers all simple signed graphs.) The same dichotomy holds if (H, Π) has no negative digons, and we conjecture that it holds always. In our proofs, we reduce s-homomorphism problems to certain ec-homomorphism problems, for which we are able to show a dichotomy. In contrast, we prove that a dichotomy theorem for ec-homomorphism problems (even when restricted to bipartite target signed graphs) would settle the dichotomy conjecture of Feder and Vardi.

1 Introduction and terminology

Graph homomorphisms and their variants play a fundamental role in the study of computational complexity. For example, the celebrated CSP Dichotomy Conjecture of Feder and Vardi [8], a major open problem in the area, can be reformulated in terms of digraph homomorphisms or graph retractions (to fixed targets). As a special case, the dichotomy theorem

*Department of Mathematics and Statistics, Thompson Rivers University, Kamloops (Canada). rbrewster@tru.ca

†Research supported by the Natural Sciences and Engineering Research Council of Canada.

‡LIMOS, Université Blaise Pascal, Clermont-Ferrand (France). florent.foucaud@gmail.com

§School of Computing Science, Simon Fraser University, Burnaby (Canada). pavol@sfu.ca

¶The author was additionally supported by the grant ERCCZ LL 1201, and part of this work was done while he was visiting the Simons Institute for the Theory of Computing.

‖CNRS - IRIF, Universit Paris Diderot, Paris (France). reza@lri.fr

of Hell and Nešetřil [15] shows that there are no NP-intermediate graph homomorphism problems. In this paper, we study homomorphisms of signed graphs from an algorithmic point of view. We study two natural types of homomorphism problems on signed graphs, one with switching and one without. For the former, we prove a dichotomy theorem for a large class of signed graphs, while for the latter we prove that a dichotomy would answer the CSP Dichotomy Conjecture in the positive.

We begin by defining signed graphs and the two types of homomorphisms. We remark that we have adopted the language of signed graphs for this paper, but readers familiar with edge-coloured graphs will recognize that our work may be equivalently formulated in terms of edge-coloured graphs. For example, the edge-coloured viewpoint is used in [25].

1.1 Signed graphs

A *signed graph* is a graph G together with a signing function $\sigma : E(G) \to \{+, -\}$. By setting $\Sigma = \sigma^{-1}(-)$, we use the notation (G, Σ) to denote this signed graph. The set Σ of negative edges is referred to as the *signature* of (G, Σ). In all our diagrams, these edges are drawn in red with dashed lines. The other edges, which are positive, are drawn in blue with solid lines. Signed graphs were introduced by Harary in [12], and studied in depth by Zaslavsky (see for example [26, 27, 28, 29, 30]). The notion of distinguishing a set Σ of edges can also be found in the work of König [18].

Signed graphs are different from 2-edge-coloured graphs with arbitrary colours, due to the fact that $\{+, -\}$ forms a group with respect to the product of signs. The most crucial of these differences comes from the following definition of the *sign* of a cycle, or more generally, a closed walk. A cycle or closed walk of (G, Σ) is said to be *negative* if the product of the signs of all the edges (considering multiplicities if an edge is traversed more than once) is the negative sign, and *positive* otherwise. A (signed) subgraph of (G, Σ) is called *balanced* if it contains no negative cycle (equivalently no negative closed walk, noting that each negative closed walk must contain a negative cycle). This notion of balance was introduced by Harary [12] and a similar idea appears in the work of König [18]. A cycle of length 2 is a *digon*. In our work, we do not consider multiple edges of the same sign, and thus we only consider negative digons.

The second notion of importance for signed graphs is the operation of *switching*, introduced by Zaslavsky [27]. To *switch* at a vertex v means to multiply the signs of all edges incident to v by $-$, that is, to switch the sign of each of these edges. (In the case of a loop at v, its sign is multiplied twice and hence it is invariant under switching.) Given signatures Σ and Σ' on a graph G, the signature Σ' is said to be *switching equivalent* to Σ, denoted $\Sigma \equiv \Sigma'$, if it can be obtained from Σ by a sequence of switchings. Equivalently, $\Sigma \equiv \Sigma'$ if their symmetric difference is an edge-cut of G.

Zaslavsky proved that two signatures Σ and Σ' of a graph G are switching equivalent if and only if they induce the same cycle signs [27]. Inherent in the proof is an algorithm to test whether Σ and Σ' are switching equivalent. For completeness, in Section 2.1, we offer an alternative certifying algorithm obtained by generalizing a method from [12].

Each of the two types of homomorphisms studied in this paper capture, in particular, the concept of proper vertex-colouring of signed graphs introduced by Zaslavsky [28] (as mappings to certain families of signed graphs).

1.2 ec-homomorphisms

Recall that given two graphs G and H, a *homomorphism* φ of G to H is a mapping of the vertices $\varphi : V(G) \to V(H)$ such that if two vertices x and y are adjacent in G, then their images are adjacent in H. We write $G \to H$ to denote the existence of a homomorphism or $\varphi : G \to H$ when we wish to explicitly name the mapping. One natural extension of this idea to signed graphs is to additionally require that homomorphisms preserve the sign of edges.

Definition 1.1. *Let (G, Σ) and (H, Π) be two signed graphs. An* ec-homomorphism *of (G, Σ) to (H, Π) is a (graph) homomorphism $\varphi : G \to H$ such that for each edge e between two vertices x and y in (G, Σ), there is an edge between $\varphi(x)$ and $\varphi(y)$ in (H, Π) having the same sign as e.*

When there exists such an ec-homomorphism, we write $(G, \Sigma) \xrightarrow{ec} (H, \Pi)$ or $\varphi : (G, \Sigma) \xrightarrow{ec} (H, \Pi)$ when we wish to explicitly name the mapping.

An ec-homomorphism $r : (G, \Sigma) \xrightarrow{ec} (H, \Pi)$ is an *ec-retraction* if (H, Π) is a subgraph of (G, Σ) and r is the identity on (H, Π). A signed graph (H, Π) is an *ec-core* if for each ec-homomorphism $\varphi : (H, \Pi) \xrightarrow{ec} (H, \Pi)$, the mapping φ is an ec-automorphism. Every signed graph (H, Π) admits an ec-retraction to a subgraph (H', Π') that is an ec-core. In fact, (H', Π') is unique up to ec-isomorphism and we call it *the ec-core* of (H, Π) [14].

The complexity of determining the existence of homomorphisms has received much attention in the literature. For classical undirected graphs, the complexity (for fixed targets) is completely determined by the dichotomy theorem of Hell and Nešetřil. Let H be a fixed graph. We define the decision problem HOM(H), also known as H-COLOURING.

HOM(H)
Instance: A graph G.
Question: Does $G \to H$?

Theorem 1.2 (Hell and Nešetřil [15])**.** *If a graph H is bipartite or contains a loop, then* HOM(H) *is polynomial-time solvable; otherwise, it is NP-complete.*

One of the questions we will consider in this work, is a possible extension of Theorem 1.2 for the class of ec-homomorphism problems for signed graphs. To this end, let (H, Π) be a fixed signed graph. We define the following decision problem.

EC-HOM(H, Π)
Instance: A signed graph (G, Σ).
Question: Does $(G, \Sigma) \xrightarrow{ec} (H, \Pi)$?

The notion of ec-homomorphisms of signed graphs (and graphs with any number of edge-colours) was studied in [1, 13, 22] from a non-computational point of view. See also [2, 3, 5, 20, 21] for studies of the computational complexity of ec-homomorphism problems.

Throughout the paper we will exploit the following (immediate) connection between graph homomorphisms and ec-homomorphisms.

Observation 1.3. *Let G and H be two graphs. Let $\varphi : V(G) \to V(H)$ be a vertex-mapping. The following are equivalent.*

(a) $\varphi : G \to H$ is a graph homomorphism.

(b) $\varphi : (G, \emptyset) \xrightarrow{ec} (H, \emptyset)$ *is an ec-homomorphism.*

(c) $\varphi : (G, E(G)) \xrightarrow{ec} (H, E(H))$ *is an ec-homomorphism.*

1.3 Constraint satisfaction problems and homomorphisms

A more general setting for the study of $\textsc{Hom}(H)$ and $\textsc{ec-Hom}(H, \Pi)$ problems is the one of general relational structures, where instead of binary relations we have relations of arbitrary arities. More formally, a relational structure S over a given *vocabulary* (a set of pairs (R_i, a_i) of relation names and arities) consists of a *domain* $V(S)$ of vertices together with a set of relations corresponding to the vocabulary, that is, $R_i \subseteq V(S)^{a_i}$ for each relation R_i of the vocabulary. Given two relational structures S and T over the same vocabulary, a homomorphism of S to T is a mapping $\varphi : V(S) \to V(T)$ such that each relation R_i is preserved; that is, for each element of R_i in S, its image in T also belongs to R_i in T. We write $S \to T$ to denote the existence of such a homomorphism.

For a fixed relational structure T (called the *template*), the *constraint satisfaction problem* for T is the decision problem defined as follows.

$\textsc{Csp}(T)$
Instance: A relational structure S over the same vocabulary as T.
Question: Does $S \to T$?

The class CSP then denotes the set of all problems of the type $\textsc{Csp}(T)$. Motivated by Theorem 1.2, Feder and Vardi [8] asked whether for every relational structure T, the problem $\textsc{Csp}(T)$ is either polynomial-time solvable or NP-complete (thus not NP-intermediate). This question has received much attention and has become known as the *Dichotomy Conjecture*.

Conjecture 1.4 (Dichotomy Conjecture, Feder and Vardi [8])**.** *For any relational structure* T, $\textsc{Csp}(T)$ *is either NP-complete or polynomial-time solvable.*

Conjecture 1.4 remains a major open problem in computational complexity. Note that the Hell-Nešetřil dichotomy theorem for graph homomorphisms [15] solves Conjecture 1.4 for templates T having just one symmetric binary relation. A number of equivalent formulations have been proposed. In particular, Feder and Vardi have shown that Conjecture 1.4 would be answered positively by a dichotomy theorem over the restricted subclass of CSP corresponding to (bipartite) digraph homomorphism problems [8], that is, templates with just one (not necessarily symmetric) binary relation.

In Section 3 of this work, we prove a similar result for ec-homomorphism problems of signed (bipartite) graphs.

1.4 s-homomorphisms

We now turn our attention to a second type of homomorphism on signed graphs (introduced by Guenin [11] and studied further by Naserasr, Rollová and Sopena in [23, 24]). These homomorphisms incorporate switching and (as we will see in Section 2.2) preserve the sign of cycles.

The notion of switching allows an extension of the theory of graph minors to signed graphs (introduced by Zaslavsky [27]) which has a stronger interplay with colouring problems. A *minor* of a signed graph is a graph obtained from a sequence of the following operations:

(i) deleting vertices or edges, (ii) contracting positive edges, and (iii) switching. Thus, the image of any negative cycle, unless it is deleted, remains negative. Furthermore, observing that negative cycles of $(G, E(G))$ are exactly the odd cycles of G, many colouring results on minor-closed graph families have been strengthened using this notion of minor. The most notable one is a strengthening of Hadwiger's conjecture proposed by Gerards and Seymour in 1980: if $(G, E(G))$ contains no $(K_n, E(K_n))$ minor, then G is $(n-1)$-colourable (this is known as the Odd Hadwiger conjecture, see [16]). Note that the case $n = 3$ of this conjecture is equivalent to claiming that all bipartite graphs are 2-colourable, whereas the original Hadwiger conjecture only asserts the 2-colourability of forests. As discussed in [24], it is then natural to study colouring through the following notion of homomorphisms of signed graphs.

Definition 1.5. *Let (G, Σ) and (H, Π) be signed graphs. An* s-homomorphism *of (G, Σ) to (H, Π) is a mapping $\varphi : V(G) \to V(H)$ such that there exist a switching (G, Σ') of (G, Σ) and a switching (H, Π') of (H, Π), such that $\varphi : (G, \Sigma') \xrightarrow{ec} (H, \Pi')$ is an ec-homomorphism.*

When there exists such an s-homomorphism, we write $(G, \Sigma) \xrightarrow{s} (H, \Pi)$, or $\varphi : (G, \Sigma) \xrightarrow{s} (H, \Pi)$ when we wish to explicitly name the mapping.

In the above definition, note that we may always assume that $\Pi' = \Pi$, since we can perform the necessary switchings on (G, Σ) instead of (H, Π). Thus we can, and will, choose a specific fixed signature for H in our proofs.

As in the case of ec-homomorphisms, an s-homomorphism of signed graphs is also a homomorphism of the underlying graphs. We prove in Section 2.2 that s-homomorphisms preserve the essential structures of signed graphs, namely adjacency and the sign of cycles. The concept of s-homomorphism was defined by Guenin in [11], where the author used this notion to capture a packing problem. Recently, the theory of s-homomorphisms was more extensively developed in [24] (see also [9, 23, 25] for subsequent studies). A related notion of homomorphisms of edge-coloured graphs where there is a switching operation is studied in [4, 20, 21]. See also [17] for homomorphisms of digraphs where a similar switching operation is allowed.

As an example, consider the vertex-mapping $\varphi : V(G) \to V(H)$ in Figure 1. As defined, φ preserves adjacency but not the sign of the edges. However, φ is an s-homomorphism of (G, Σ) to (H, Π). Indeed, there exists a switching of Σ to Σ' (defined in Figure 1) such that φ preserves edges and their signs. In fact, several switchings of (G, Σ) exist such that the given vertex-mapping preserves edges and their signs.

As for ec-homomorphisms, a signed graph (H, Π) is defined to be an *s-core* if for every s-homomorphism $\varphi : (H, \Pi) \xrightarrow{s} (H, \Pi)$, φ is an s-automorphism. *The* s-core of (H, Π) is defined as before, indeed it is again easy to prove that every signed graph (H, Π) has an s-core that is unique up to s-isomorphism and switching.

Let (H, Π) be a fixed signed graph. We define the s-homomorphism decision problem for (H, Π) analogously as in the ec-case.

s-Hom(H, Π)
Instance: A signed graph (G, Σ).
Question: Does $(G, \Sigma) \xrightarrow{s} (H, \Pi)$?

A primary concern in this work is the computational complexity of the s-Hom(H, Π) problem. By Observation 1.3, we note that when Π is switching equivalent to $E(H)$ or to $\emptyset$, then s-Hom(H, Π) has the same complexity as Hom(H). To see this, given an input signed graph (G, Σ), we can decide in polynomial time whether Σ is switching equivalent to

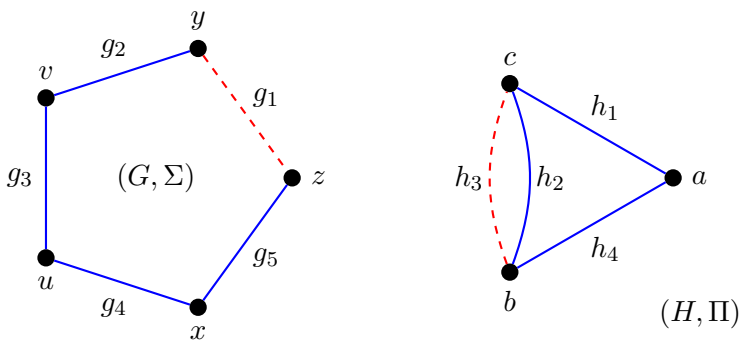

Figure 1: The mapping $\varphi : V(G) \to V(H)$ is defined by $\varphi(z) = a, \varphi(y) = \varphi(u) = b, \varphi(v) = \varphi(x) = c$. After switching at y and v, this is an edge and sign-preserving mapping. Switching at y and u is a second possibility.

$E(G)$ or to $\emptyset$ (see Proposition 2.1). Then, the computational complexity of the problem is decided by Theorem 1.2. Prior to this work, the only other known case of the problem's complexity is the study of [9] about signed cycles, where it is proved that s-Hom(C_{2k}, Π) is NP-complete if Π has an odd number of elements, and polynomial-time solvable otherwise (note that s-Hom(C_{2k+1}, Π) for $k \geq 1$ is always NP-complete). Here, we give a full dichotomy characterization in the case where (H, Π) is a simple signed graph. Indeed we prove an even stronger result.

Theorem 4.2. *Let (H, Π) be a connected signed graph that does not contain all three of a negative digon, a negative loop, and a positive loop. Then,* s-Hom(H, Π) *is polynomial-time solvable if the s-core of (H, Π) has at most two edges; it is NP-complete otherwise.*

We believe that a full dichotomy holds for s-homomorphism problems. In fact, we conjecture that all cases not covered by Theorem 4.2 are NP-complete.

Conjecture 1.6. *Let (H, Π) be a connected signed graph. Then,* s-Hom(H, Π) *is polynomial-time solvable if the s-core of (H, Π) has at most two edges; it is NP-complete otherwise.*

Note that the polynomial half of Conjecture 1.6 is proved in the proof of Theorem 4.2, that is, the polynomial case holds for all signed graphs.

1.5 Structure of the paper

The paper is organized as follows. In Section 2, we present preliminary results on both kinds of signed graph homomorphisms, some of which are fundamental for our subsequent proofs. In particular, we show that the problem of testing the existence of an s-homomorphism between signed graphs can be captured through a special construction, *switching graphs*, using ec-homomorphisms. In Section 3, we prove the equivalence between a dichotomy for ec-homomorphism problems for bipartite signed graphs and a dichotomy for all of CSP. We then prove our main result, Theorem 4.2, in Section 4. Finally, in Section 5, we conclude with some remarks and questions.

2 Preliminaries

In this section, we give our algorithm for testing switching equivalence of two signatures, provide an equivalent definition for s-homomorphisms, introduce the switching graph construction, and connect our work to vertex-colourings of signed graphs.

2.1 A certifying algorithm for switching equivalence

The proof of the following proposition is a certifying algorithm obtained by generalizing a method from [12].

Proposition 2.1. *Given two signed graphs (G, Σ) and (G, Σ'), it can be decided in polynomial time whether $\Sigma \equiv \Sigma'$.*

Proof. Let $B = \Sigma \triangle \Sigma'$. We wish to test if there is an edge-cut $(X, \overline{X})$ such that $B = E(X, \overline{X})$. We present a certifying algorithm that either finds such a cut or returns a cycle whose sign is different in (G, Σ) and (G, Σ'). Let $W_1, W_2, \ldots, W_t$ be the components of $G \setminus B$. As all edges of each W_i have the same sign in both Σ and Σ', each component W_i must be entirely contained by one side of the desired edge-cut. Thus, the aim is to partition the components into two sets X and Y so that the edges of B have one end in X and the other in Y. In other words, is it true that the graph resulting from contracting each W_i to a single vertex is bipartite? If yes, then the edges of the resulting bipartite graph form the cut B. Switching on this cut transforms Σ to Σ'. If no, then there is a cycle C in G containing an odd number of edges from B. This cycle is positive in exactly one of Σ or Σ' (and negative in the other), certifying that $\Sigma \not\equiv \Sigma'$. □

In the case that Σ and Σ' are not switching equivalent, the certifying cycle C yields the following corollary due to Zaslavsky (as mentioned in the introduction).

Corollary 2.2 (Zaslavsky [27])**.** *Two signatures Σ and Σ' of the same graph are switching equivalent if and only if they induce the same cycle signs.*

2.2 An equivalent definition of s-homomorphisms

Both homomorphisms of classical graphs and ec-homomorphisms of signed graphs are vertex-mappings that preserve adjacency (and the adjacency signs for the latter). Such a homomorphism φ naturally defines the associated mapping of the edges (see the book [14, page 4]).

In the case of s-homomorphisms, the edge-mapping $\varphi^{\#}$ may not be well-defined by the images of the vertices alone (when there are negative digons or a vertices with two loops of opposite signs), as $\varphi^{\#}$ may depend on the choice of switching. Recall the example of Figure 1. We had several possible switchings of (G, Σ) so that $\varphi : V(G) \to V(H)$ induces a sign-preserving mapping, and therefore $\varphi^{\#}$ depends on both the vertex-mapping and the particular switching of (G, Σ) one has chosen. Hence, we will, when required, explicitly specify the associated mapping $\varphi^{\#}$.

Next, we show that one may view an s-homomorphism as a vertex-mapping φ with an associated edge-mapping $\varphi^{\#}$ that preserves the important structures in signed graphs: the adjacency of vertices and the signs of cycles.

Theorem 2.3. *Let (G,Σ) and (H,Π) be signed graphs. Then $\varphi : V(G) \to V(H)$, with a given associated mapping $\varphi^{\#} : E(G) \to E(H)$, is an s-homomorphism $(G,\Sigma) \xrightarrow{s} (H,\Pi)$ if and only if $\varphi : G \to H$ is a homomorphism of the underlying graphs and, for every cycle C of G, the sign of C in (G,Σ) is the same as the sign of the closed walk $\varphi^{\#}(C)$ in (H,Π).*

Proof. First assume that φ is an s-homomorphism of (G,Σ) to (H,Π). Then, by the definition, φ is an ec-homomorphism of (G,Σ') to (H,Π) for some signature Σ' that is switching equivalent to Σ. Clearly, the sign of C in (G,Σ') is the same as the sign of $\varphi^{\#}(C)$ in (H,Π), but it is also the same as the sign of C in (G,Σ) since Σ and Σ' are switching equivalent.

Now, suppose that the condition holds for each cycle. Let Σ' be the inverse image of Π by $\varphi^{\#}$. Our claim is that Σ' and Σ are switching equivalent; this would prove that φ is an s-homomorphism of (G,Σ) to (H,Π). To prove our claim, using Corollary 2.2, it is enough to show that each cycle has the same sign in (G,Σ) and (G,Σ'). But this is indeed the case, as they both are the same as the sign of $\varphi^{\#}(C)$ in (H,Π), one by the condition of the theorem, the other from the definition of Σ'. □

Returning to the example in Figure 1, the vertex-mapping g with $g(u) = b, g(v) = g(z) = c, g(x) = g(y) = a$ maps only one edge g_3 to the negative digon. In light of Theorem 2.3, we must have $g^{\#}(g_3) = h_3$.

2.3 Switching graphs

We now describe a construction that is crucial in our proofs.

Definition 2.4. *Let (G,Σ) be a signed graph. The* switching graph *of (G,Σ) is a signed graph denoted $P(G,\Sigma)$ and constructed as follows.*

(i) For each vertex u in $V(G)$ we have two vertices u_0 and u_1 in $P(G,\Sigma)$.

(ii) For each edge e between u and v in G, we have four edges between u_i and v_j $(i,j \in \{0,1\})$ in $P(G,\Sigma)$, with the edges between u_i and v_i having the same sign as e and the edges between u_i and v_{1-i} having the opposite sign as e $(i \in \{0,1\})$. (In particular, loops do not change sign.)

See Figures 2 and 4 for examples of signed graphs and their switching graphs. The notion of switching graph was defined by Brewster and Graves in [4] in a more general setting related to permutations (they used the term *permutation graph*). Their work built on that of Klostermeyer and MacGillivray [17] who used a similar definition in the context of digraphs. The construction is also used in [25]. Zaslavsky used a construction similar (but different) to that of switching graphs [28].

Let (G,Σ) be a signed graph and (G,Σ') be any switching equivalent signed graph, that is, $\Sigma \equiv \Sigma'$. A fundamental property of $P(G,\Sigma)$ is that it contains as a subgraph both (G,Σ) and (G,Σ'). That is, it contains as a subgraph all signed graphs that are switching equivalent to (G,Σ). The following proposition allows us to transform questions about s-homomorphisms to the setting of ec-homomorphisms.

Proposition 2.5. *Let (G,Σ) and (H,Π) be two signed graphs. The following are equivalent.*

(a) $(G,\Sigma) \xrightarrow{s} (H,\Pi)$,

(b) $(G,\Sigma) \xrightarrow{ec} P(H,\Pi)$,

(c) $P(G,\Sigma) \xrightarrow{ec} P(H,\Pi)$.

Proof. (a) $\iff$ (b): Let the vertices of H be $\{v_1, v_2, \ldots, v_n\}$ and let the corresponding paired vertices in $P(H,\Pi)$ be $v_{i,0}$ and $v_{i,1}$ for $i = 1, 2, \ldots, n$. Suppose $(G,\Sigma) \xrightarrow{s} (H,\Pi)$. Thus, there is a signed graph (G,Σ') with $\Sigma' \equiv \Sigma$ and an ec-homomorphism $\varphi : (G,\Sigma') \xrightarrow{ec} (H,\Pi)$.

Let $E(X,\overline{X}) = \Sigma \triangle \Sigma'$ be the edge-cut certifying the switching equivalence of the two signatures. Define $\psi : (G,\Sigma) \xrightarrow{ec} P(H,\Pi)$ by

$$\psi(u) = \begin{cases} v_{i,0} & \text{if } \varphi(u) = v_i \text{ and } u \in X \\ v_{i,1} & \text{if } \varphi(u) = v_i \text{ and } u \in \overline{X} \end{cases}.$$

It is straightforward to verify that ψ is an ec-homomorphism.

For the converse, we have $P(H,\Pi) \xrightarrow{s} (H,\Pi)$ (switching on the edges between the two copies of (H,Π) and projecting $v_{i,j} \mapsto v_i$ is an s-homomorphism).

(b) $\iff$ (c): First observe that $(G,\Sigma) \subseteq P(G,\Sigma)$. Thus, $P(G,\Sigma) \xrightarrow{ec} P(H,\Pi)$ implies $(G,\Sigma) \xrightarrow{ec} P(H,\Pi)$. On the other hand, suppose that $\varphi : (G,\Sigma) \xrightarrow{ec} P(H,\Pi)$ is an ec-homomorphism. We construct an ec-homomorphism $\psi : P(G,\Sigma) \xrightarrow{ec} P(H,\Pi)$ as follows: if $\varphi(u) = v_{i,j}$, then $\psi(u_0) = v_{i,j}$ and $\psi(u_1) = v_{i,1-j}$. □

As s-cores of signed graphs are our fundamental object of study, we remark that Theorem 15 of [4] characterizes the s-cores of signed graphs in terms of switching graphs. (There is a technical requirement we must make on the $P(G,\Sigma)$ construction in order to apply the following theorem. If for some vertex v of G, v_0 and v_1 have the same sets of positive and negative neighbours in $P(G,\Sigma)$, then we delete vertex v_1 from $P(G,\Sigma)$. This corresponds to the case where the multiset of colours of incident edges of a loop-free vertex v are invariant under switching, that is, v is an isolated vertex, or v is incident only with negative digons.)

Theorem 2.6 (Brewster and Graves [4, Theorem 15])**.** *Let (G,Σ) be a signed graph. Then, (G,Σ) is an s-core if and only if $P(G,\Sigma)$ is an ec-core.*

2.4 Signed graph vertex-colourings

Zaslavsky introduced and studied vertex-colourings of signed graphs in a series of papers [28, 29, 30]. In this section, we formulate these colourings in the language of homomorphisms of signed graphs to particular targets. These targets play the same role (for signed graph colourings) that complete graphs play for colourings of classical graphs. (Recall that a homomorphism of a graph to the complete graph K_k corresponds to a proper k-vertex-colouring.)

In fact, Zaslavsky defined two types of colouring. Let k be a positive integer. A *proper k-colouring* of a signed graph (G,Σ) is a vertex-mapping $\phi : V(G) \to \{0, \pm 1, \ldots, \pm k\}$ with the property that for any two adjacent vertices u and v of G, $\phi(u) \cdot \sigma(e) \neq \phi(v)$ where e is an edge incident with u and v and $\sigma(e)$ is its sign. Thus, two vertices joined by a positive edge cannot receive the same colour, while two vertices joined by a negative edge cannot receive opposite colours. Further, the colouring is *zero-free* if it maps no vertex to 0.

We now formulate signed graph colourings as homomorphisms. Let k be a positive integer. The signed graph (Z_k, Υ_k) has vertex set $\{0, 1, \ldots, k\}$ and edges consisting of a negative digon between all pairs of distinct vertices and negative loops on $\{1, 2, \ldots, k\}$. The signed graph

(Z_k^*, Υ_k^*) is obtained from (Z_k, Υ_k) by deleting 0 and its incident edges. See Figure 2 for examples. The following proposition is immediate using Proposition 2.5.

Proposition 2.7. *Let (G, Σ) be a signed graph. The following three statements are equivalent.*

(a) (G, Σ) *admits a proper k-colouring.*

(b) $(G, \Sigma) \xrightarrow{s} (Z_k, \Upsilon_k)$.

(c) $(G, \Sigma) \xrightarrow{ec} P(Z_k, \Upsilon_k)$.

In addition, the following three statements are equivalent.

(d) (G, Σ) *admits a proper zero-free k-colouring.*

(e) $(G, \Sigma) \xrightarrow{s} (Z_k^*, \Upsilon_k^*)$.

(f) $(G, \Sigma) \xrightarrow{ec} P(Z_k^*, \Upsilon_k^*)$.

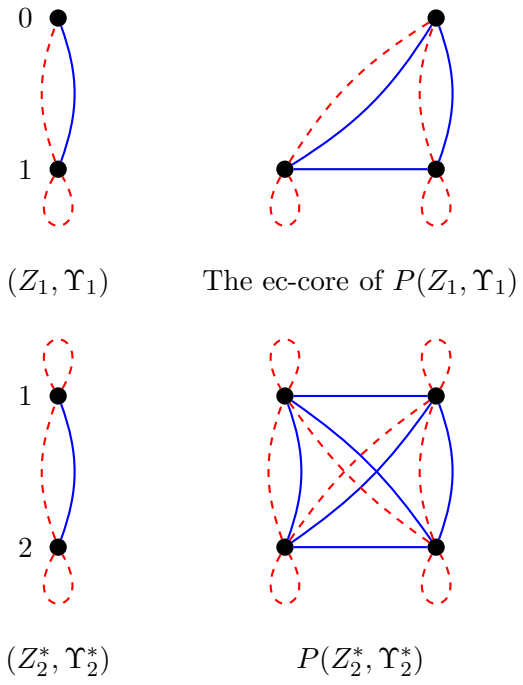

Figure 2: The graphs (Z_1, Υ_1), (Z_2^*, Υ_2^*) and (the ec-cores of) their switching graphs.

Since (Z_k, Υ_k) and (Z_k^*, Υ_k^*) do not contain positive loops, Theorem 4.2 is powerful enough to classify the complexity of signed graph colourings. Specifically, for all positive k, the s-core of (Z_k, Υ_k) has at least three edges and thus the corresponding colouring problem is NP-complete. For zero-free colourings, the s-core of (Z_k^*, Υ_k^*) consists of a single negative loop when $k = 1$, and at least three edges for $k \geq 2$. Hence, the corresponding colouring problem is polynomial-time in the former case and NP-complete in the latter.

3 A dichotomy for ec-homomorphisms of bipartite signed graphs implies a dichotomy for all of CSP

We now show the equivalence between the set of EC-HOM(H, Π) problems where (H, Π) is a signed graph, and all of CSP. We will require the two following retraction decision problems in the reduction proved below. Formally, let H be a fixed graph and (H, Π), a fixed signed graph.

RET(H)
Instance: A graph G containing H as a subgraph.
Question: Is there a retraction of G to H?

EC-RET(H, Π)
Instance: A signed graph (G, Σ) containing (H, Π) as a subgraph.
Question: Is there a retraction $r : (G, \Sigma) \xrightarrow{ec} (H, \Pi)$?

An *alternating path* is a (signed) path whose edges alternate negative and positive. Following the construction of Feder and Vardi (Theorem 10 of [8]), and the development of these ideas in the book [14, Chapter 5.3], we have the following theorem which shows that a dichotomy theorem for the set of EC-HOM(H, Π) problems would provide a positive answer to the Feder and Vardi Dichotomy Conjecture.

Theorem 3.1. *For each CSP template T, there is a signed graph (H, Π) such that* EC-HOM(H, Π) *and* CSP(T) *are polynomially equivalent. Moreover, (H, Π) can be chosen to be bipartite and homomorphic to an alternating path.*

Proof. We follow the proof of Theorem 5.14 in the book [14] proving a similar statement for digraph homomorphism problems. The structure of the proof in [14] is as follows. First, one shows that for each CSP template T, there is a bipartite graph H such that the CSP(T) problem and the RET(H) problem are polynomially equivalent. Next, it is shown that for each bipartite graph H, there is a digraph H' such that RET(H) and RET(H') are polynomially equivalent. Finally, it is observed that H' is a core and thus RET(H') and HOM(H') are polynomially equivalent. We will adapt this proof to the case of ec-homomorphism problems of signed graphs.

The construction of H' from H in [14] is through the use of so-called zig-zag paths. For signed graphs, we construct a similar collection of paths. This will allow us to construct a signed graph (H', Π) from a bipartite graph H such that RET(H) and EC-RET(H', Π) are polynomially equivalent. Our paths will have positive edges denoted by B and negative edges denoted by R. Hence, the path BR^3B^3R consists of one positive edge, three negative edges, three positive edges and a negative edge. The maximal monochromatic (constant signs) subpaths are called *runs*. Thus, the above path is the concatenation of four runs: the first and last of length 1, and the middle two of length 3.

Given an odd integer ℓ, we construct a signed path P consisting of ℓ runs. The first and the last run each consist of a single negative edge. The interior runs are of length 3. We denote that last (rightmost) vertex of P by 0. From P, we construct $\ell - 2$ paths $P_1, \ldots, P_{\ell-2}$. Path P_i $(i = 1, 2, \ldots, \ell - 2)$ is obtained from P by replacing the i^{th} run of length 3 with a run of length 1. We denote the rightmost vertex of P_i by i.

Similarly, for an even integer k, we construct a second family of paths Q and Q_j, ($j = 1, 2, \ldots, k-2$). The leftmost vertex of Q is 1 and the leftmost vertex of Q_j is j. The paths are described below:

$$
\begin{array}{rclcrcl}
P & := & R\underbrace{B^3R^3\cdots R^3B^3}_{\ell-2}R & & Q & := & R\underbrace{B^3R^3\cdots B^3R^3}_{k-2}B \\
P_i & := & R\underbrace{B^3\cdots R^3}_{i-1}B\underbrace{R^3\cdots B^3}_{\ell-i-2}R \quad (i \text{ odd}) & & Q_j & := & R\underbrace{B^3\cdots R^3}_{j-1}B\underbrace{R^3\cdots R^3}_{k-j-2}B \quad (j \text{ odd}) \\
P_i & := & R\underbrace{B^3\cdots B^3}_{i-1}R\underbrace{B^3\cdots B^3}_{\ell-i-2}R \quad (i \text{ even}) & & Q_j & := & R\underbrace{B^3\cdots B^3}_{j-1}R\underbrace{B^3\cdots R^3}_{k-j-2}B \quad (j \text{ even})
\end{array}
$$

We observe the following (see [14, page 156]):

1. The paths P and P_i ($i = 1, 2, \ldots \ell - 2$) each admit an ec-homomorphism *onto* an *alternating path* of length ℓ, (that is, a path consisting of ℓ runs each of length one: $RBRB\cdots R$).
2. The paths Q and Q_j ($j = 1, 2, \ldots k - 2$) each admit an ec-homomorphism onto an alternating path of length k.
3. $P_i \xrightarrow{ec} P_{i'}$ implies $i = i'$.
4. $Q_j \xrightarrow{ec} Q_{j'}$ implies $j = j'$.
5. $P \xrightarrow{ec} P_i$ for all i.
6. $Q \xrightarrow{ec} Q_j$ for all j.
7. if X is a signed graph and x is a vertex of X such that $f : X \xrightarrow{ec} P_i$ and $f' : X \xrightarrow{ec} P_{i'}$ for $i \neq i'$ with $f(x) = i$ and $f'(x) = i'$, then there is an ec-homomorphism $F : X \xrightarrow{ec} P$ with $F(x) = 0$.
8. if Y is a signed graph and y is a vertex of Y such that $f : Y \xrightarrow{ec} Q_j$ and $f' : Y \xrightarrow{ec} Q_{j'}$ for $j \neq j'$ with $f(y) = j$ and $f'(y) = j'$, then there is an ec-homomorphism $F : Y \xrightarrow{ec} Q$ with $F(y) = 1$.

We note that alternating paths in signed graphs can be used to define *height* analogously to height in directed acyclic graphs. Specifically, suppose (G, Σ) is a connected signed graph that admits an ec-homomorphism onto an alternating path, say AP. Let the vertices of AP be $h_0, h_1, \ldots, h_t$. Observe that each vertex in the path has at most one neighbour joined by a negative edge and at most one neighbour joined by a positive edge. Thus, once a single vertex u in (G, Σ) is mapped to AP, the image of each neighbour of u is uniquely determined; by connectivity, the image of all vertices is uniquely determined. In particular, as (G, Σ) maps *onto* AP, there is exactly one ec-homomorphism of (G, Σ) to AP. (More precisely, if the path has odd length, there is an ec-automorphism that reverses the path. In this case there are two ec-homomorphisms that are equivalent up to reversal.) We then observe that if $g : (G, \Sigma) \xrightarrow{ec} AP$ with g being onto, $h : (H, \Pi) \xrightarrow{ec} AP$, and $f : (G, \Sigma) \xrightarrow{ec} (H, \Pi)$, then for all vertices $u \in V(G)$, $g(u) = h(f(u))$. This allows us to define the height of $u \in V(G)$ to be

h_i when $g(u) = h_i$. Specifically, vertices at height h_i in (G, Σ) must map to vertices at height h_i in (H, Π).

For each problem T in CSP, there is a bipartite graph H such that $\textsc{Csp}(T)$ and $\textsc{Ret}(H)$ are equivalent [8, 14]. Let H be a bipartite graph with bipartition (A, B), where $A = \{a_1, \ldots, a_{|A|}\}$ and $B = \{b_1, \ldots, b_{|B|}\}$. Let ℓ (respectively k) be the smallest odd (respectively even) integer greater than or equal to $|A|$ (respectively $|B|$). Recall that H has as its core a single edge, for which the $\textsc{Hom}(H)$ problem is polynomial. However, we are using the $\textsc{Ret}(H)$ problem and thus we require the retraction to be the identity on H. We will now add some signed gadgets to the vertices of H to enforce that any homomorphism from the copy of H in G to H itself must act as the identity on H. To each vertex $a_i \in A$, attach a copy of P_i identifying i in P_i with a_i in A. To each vertex $b_j \in B$ attach a copy of Q_j, identifying j in Q_j with b_j in B. Let all original edges of H be positive. Call the resulting signed graph (H', Π). See Figure 3 for an illustration.

Let G be an instance of $\textsc{Ret}(H)$. In particular, we may assume without loss of generality that H is a subgraph of G, G is connected, and G is bipartite. Let (A', B') be the bipartition of G where $A \subseteq A'$ and $B \subseteq B'$. To each vertex v of $A' \setminus A$, we attach a copy of P, identifying v and 0. To the vertices of $A \cup B$, we attach paths P_i and Q_j as described above to create a copy of H'. We let the original edges of G be positive. Call the resulting signed graph (G', Σ). In particular, note that (G', Σ) and (H', Π) both map onto an alternating path of length $\ell + k + 1$. The (original) vertices of G and H are at height ℓ and $\ell + 1$ for parts A and B respectively. In particular, by the eight above properties, under any ec-homomorphism $f : (G', \Sigma) \xrightarrow{ec} (H', \Pi)$ the restriction of f to G must map onto H with vertices in A' mapping to A and vertices in B' mapping to B.

Using the eight properties of the paths above and following the proof of Theorem 5.14 in [14], we conclude that G is a YES instance of $\textsc{Ret}(H)$ if and only if (G', Σ) is a YES instance of $\textsc{ec-Ret}(H', \Pi)$.

On the other hand, let (G', Σ) be an instance of $\textsc{Ret}(H', \Pi)$. We sketch the proof from [14]. We observe that (G', Σ) must map to an alternating path of length $\ell + k + 1$. The two levels of (G', Σ) corresponding to H induce a bipartite graph (with positive edges) which we call G. The components of the subgraph of (G', Σ) obtained by removing the edges of G fall into two types: those which map to lower levels and those which map to higher levels than G. Let C_t be a component that maps to a lower level. After required identifications, we may assume that C_t contains only one vertex from G (say v) and C_t must map to some P_i. If P_i is the unique P_i path to which C_t maps, then we modify G' by identifying v and i. Otherwise, C_t maps to two paths and hence to all paths. The resulting signed graph has an ec-retraction to (H', Π) if and only if G has a retraction to H. □

4 A dichotomy theorem for s-homomorphism problems

In this section, we prove a dichotomy theorem for s-homomorphism problems $\textsc{s-Hom}(H, \Pi)$ such that (H, Π) belongs to the family $\mathcal{S}^*$ consisting of those signed graphs not containing all three of a negative digon, a positive and a negative loop. This includes a full dichotomy for all simple signed graphs. We remark that, by Proposition 2.5, our theorem gives a dichotomy theorem for ec-homomorphism problems $\textsc{ec-Hom}(P(H, \Pi))$ where (H, Π) (equivalently, $P(H, \Pi)$) belongs to $\mathcal{S}^*$. Recall that by Theorem 3.1, a full dichotomy for ec-homomorphism problems

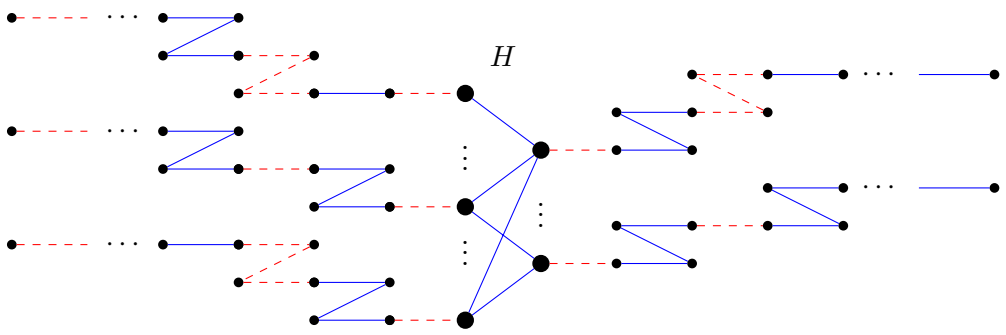

Figure 3: Construction of a signed target (H', Π) from a RET(H) problem.

for all signed bipartite graphs would imply a dichotomy theorem for all of CSP, and thus settle the Dichotomy Conjecture of Feder and Vardi.

4.1 The indicator construction

We recall the *indicator construction* defined in [15]. Given our setting, we use the construction for signed graphs, but it can be generalized to any number of edge-colours and any number of simultaneous indicators (see [2]).

Let (H, Π) be a signed graph. An *indicator*, $((I, \Lambda), i, j)$, is a signed graph (I, Λ) with two distinguished vertices i and j such that (I, Λ) admits an ec-automorphism mapping i to j and vice-versa. The *result of the indicator* $((I, \Lambda), i, j)$ applied to (H, Π) is an undirected graph denoted $(H, \Pi)^*$ and defined as follows.

(i) $V((H, \Pi)^*) = V(H)$

(ii) There is an edge from u to v in $(H, \Pi)^*$ if there is an ec-homomorphism of $(I, \Lambda) \xrightarrow{ec} (H, \Pi)$ such that $i \mapsto u$ and $j \mapsto v$.

Note that $(H, \Pi)^*$ is a classical undirected graph. The notation reflects the fact that it is derived from the signed graph (H, Π). The indicator construction is a fundamental tool in proving NP-completeness results.

Theorem 4.1 (Hell and Nešetřil [15])**.** *The* HOM$((H, \Pi)^*)$ *problem admits a polynomial-time reduction to* EC-HOM(H, Π)*. In particular, if* HOM$((H, \Pi)^*)$ *is NP-complete, then* EC-HOM(H, Π) *is NP-complete.*

We remark that if one removes the requirement that $((I, \Lambda), i, j)$ admits an ec-automorphism interchanging i and j, Theorem 4.1 still holds, but then $(H, \Pi)^*$ has directed edges.

As an example, consider the signed graph (H, Π) and the switching graph $P(H, \Pi)$ in Figure 4. Let (I, Λ) be a path of length 2 with endpoints i, j and middle vertex c, with ic positive and cj negative. The result of the indicator $((I, \Lambda), i, j)$ on $P(H, \Pi)$ is a directed graph $P(H, \Pi)^*$. However, one can easily verify that in $P(H, \Pi)$, if there exists $\varphi : (I, \Lambda) \xrightarrow{ec} P(H, \Pi)$ with $\varphi(i) = u$ and $\varphi(j) = v$, then there exists $\varphi' : (I, \Lambda) \xrightarrow{ec} P(H, \Pi)$ such that $\varphi'(i) = v$ and $\varphi'(j) = u$. (Specifically, if $\varphi(c) = w_t$, then use $\varphi'(c) = w_{1-t}$.) Thus, $P(H, \Pi)^*$

is a symmetric digraph which we may view as an undirected graph. Alternatively, take two copies of (I, Λ) and identify i in the first with j in the second, and vice versa. The result is an indicator $((I', \Lambda'), i, j)$ consisting of a 4-cycle whose edges alternate positive and negative, where i and j are antipodal vertices. Now, the result of the indicator construction (with respect to $((I', \Lambda'), i, j)$) is the undirected graph $P(H, \Pi)^*$. See Figure 4.

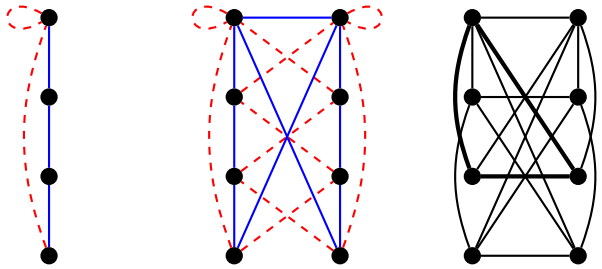

Figure 4: A signed graph (G, Σ), the switching graph $P(G, \Sigma)$, and the result of the indicator construction $P(G, \Sigma)^*$.

In the proofs below, we will use the indicator $((I, \Lambda), i, j)$ defined above for ease of explanation with the understanding that the resulting graph $P(H, \Pi)^*$ is an undirected graph (for either of the reasons above).

4.2 The dichotomy theorem

We now prove our main theorem.

Theorem 4.2. *Let (H, Π) be a connected signed graph that does not contain all three of a negative digon, a positive loop, and a negative loop. Then, s-Hom(H, Π) is polynomial-time solvable if the s-core of (H, Π) has at most two edges; it is NP-complete otherwise.*

Proof. Suppose that (H, Π) is a connected signed graph that does not contain all three of a negative digon, a positive loop, and a negative loop.

Polynomial cases.

Suppose that the s-core of (H, Π) has at most two edges and at least one is a loop. It is straightforward to check that the s-core consists of either: a single vertex with a loop of each sign, or a single vertex with a single loop.

If the s-core of (H, Π) is a single vertex with both kinds of loops, then every signed graph trivially maps to this graph. Thus, for the remainder of the proof, particularly the NP-complete cases below, (H, Π) will not have a vertex with both kinds of loops.

If the s-core of (H, Π) is a single vertex with a single loop, $(G, \Sigma) \xrightarrow{s} (H, \Pi)$ if and only if $\Sigma \equiv E(G)$ in the case where the loop is negative or $\Sigma \equiv \emptyset$ in the case where the loop is positive (that is, (G, Σ) can be switched so that all edges have the same sign as the loop in the s-core of (H, Π)); this can be checked in polynomial time by Proposition 2.1.

Thus, assume that the s-core of (H,Π) is loop-free. First, suppose that there is a single edge joining two vertices. (Note that a single positive edge is switching equivalent to a single negative edge.) This is the case if and only if H is bipartite and all cycles of (H,Π) are positive, that is, $\Pi \equiv \emptyset$. Then, $(G,\Sigma) \xrightarrow{s} (H,\Pi)$ if and only if G is loop-free and bipartite, and all cycles of (G,Σ) are positive. All conditions are easy to check (using Proposition 2.1 for the third one).

If the s-core of (H,Π) consists of a single negative digon, $(G,\Sigma) \xrightarrow{s} (H,\Pi)$ if and only if G is loop-free and bipartite. This condition can again be checked in polynomial time.

There is no other signed graph with at most two edges that is an s-core.

NP-complete cases.

Assume that the s-core of (H,Π) has at least three edges. We consider several cases.

Case (a): (H,Π) *contains a negative digon.*

If H is loop-free, then H must contain an odd cycle, as otherwise the negative digon is the s-core. Then, $P(H,\Pi)$ contains an odd cycle of the same length and sign, positive or negative. Restricting the input of EC-HOM$(P(H,\Pi))$ to graphs with only positive (respectively negative) edges shows that EC-HOM$(P(H,\Pi))$ is NP-complete by the Hell-Nešetřil dichotomy of Theorem 1.2. Thus, by Proposition 2.5, S-HOM(H,Π) is NP-complete.

On the other hand, suppose that (H,Π) contains a loop. By the assumptions about (H,Π) in the statement of the theorem, all loops have the same sign. Let us assume that they are all negative (the proof is symmetric if they are all positive). Let u,v be the vertices of a negative digon of (H,Π), and let w be a vertex with a loop. Since H is connected, there is a path P from v to w. Possibly after interchanging the roles of u and v, we may assume without loss of generality that $P = v,u,x_1,x_2,\ldots,w$ where $u = w$ is allowed. Moreover, by switching, we may assume that each edge of P is positive. Let the resulting signature be Π'. Then, in $P(H,\Pi')$, the cycle $v_0,u_0,x_{1,0},x_{2,0},\ldots,w_0,w_1,\ldots,x_{2,1},x_{1,1},u_1,v_0$ is an odd cycle consisting of only positive edges. By assumption, (H,Π) has no positive loops. Thus, EC-HOM$(P(H\Pi'))$ restricted to instances with only positive edges is NP-complete by the Hell-Nešetřil dichotomy of Theorem 1.2 and S-HOM(H,Π) is NP-complete by Proposition 2.5. This settles Case (a).

Thus, for the remainder of the proof, assume that (H,Π) has no negative digon.

Case (b): (H,Π) *contains a negative even cycle* C^-_{2k} *(*$k \geq 2$*).*

Without loss of generality, we may assume (after appropriate switching) that in (H,Π), the cycle C^-_{2k} contains $2k-1$ positive edges, say $v_1v_2, v_2v_3, \ldots, v_{2k-1}v_{2k}$, and a single negative edge $v_{2k}v_1$. Construct the graph $P(H,\Pi)$. Let the paired vertices be $v_{i,0}$ and $v_{i,1}$ for $i \in \{1,\ldots,2k\}$. Thus $v_{1,0},v_{2,0},\ldots,v_{2k,0},v_{1,0}$ and $v_{1,1},v_{2,1},\ldots,v_{2k,1},v_{1,1}$ are two copies of C^-_{2k}.

Let (I,Λ) be a path of length 2 on vertices $\{i,c,j\}$ with a positive edge ic and a negative edge cj. Let $P(H,\Pi)^*$ be the result of the indicator construction on $P(H,\Pi)$ with respect to $((I,\Lambda),i,j)$. From the comments above, we recall that $P(H,\Pi)^*$ is an undirected graph. We will show that $P(H,\Pi)^*$ is loop-free and contains an odd cycle. Theorem 1.2 then implies that HOM$(P(H,\Pi)^*)$ is NP-complete, from which it follows that EC-HOM$(P(H,\Pi))$ is NP-complete by Theorem 4.1, and thus S-HOM(H,Π) is NP-complete by Proposition 2.5.

To see that $P(H,\Pi)^*$ is loop-free, observe that an ec-homomorphic image $f(I,\Lambda)$ of (I,Λ) where $f(i) = f(j)$ must be a negative digon or a pair of loops (one positive, one negative at

a single vertex). Since $P(H,\Pi)$ does not contain either structure, $P(H,\Pi)^*$ is loop-free. Next, we observe that there is a path $v_{1,0}, v_{2,0}, v_{3,1}$ that is a copy of (I,Λ) in $P(H,\Pi)$. Hence, there is an edge from $v_{1,0}$ to $v_{3,1}$ in $P(H,\Pi)^*$. There is also a copy of (I,Λ) on $v_{3,1}, v_{2,1}, v_{3,0}$, giving a path in $P(H,\Pi)^*$: $v_{1,0}, v_{3,1}, v_{3,0}$. Continuing, we obtain a path of even length $v_{1,0}, v_{3,1}, v_{3,0}, v_{5,1}, v_{5,0}, \ldots, v_{2k-1,1}, v_{2k-1,0}$. Finally, the path $v_{2k-1,0}, v_{2k,0}, v_{1,0}$ is a copy of (I,Λ) in $P(H,\Pi)$. Hence, there is an edge from $v_{2k-1,0}$ to $v_{1,0}$ in $P(H,\Pi)^*$. This implies the existence of the required odd cycle in $P(H,\Pi)^*$, completing Case (b). (See for example Figure 4.)

We now assume that (H,Π) has no negative digon and all even cycles are positive. In particular, $\Pi \equiv \emptyset$ if and only if all *odd* cycles (including loops) are positive. Similarly, $\Pi \equiv E(H)$ if and only if all odd cycles are negative.

Case (c): (H,Π) has an odd cycle C of length at least 3, but no loop with the same sign as C.

Let us assume that C is positive and there is no positive loop. (The case where C is negative and there is no negative loop can be handled symmetrically.) In $P(H,\Pi)$, there is a copy of C (with positive edges), but no positive loop. Hence, as in Case (a), EC-HOM$(P(H,\Pi))$ is NP-complete and thus, by Proposition 2.5, S-HOM(H,Π) is NP-complete, settling Case (c).

Case (d): For each odd cycle C in (H,Π), there is a loop of the same sign as C.

If (H,Π) is loop-free, then we conclude that it has no odd cycles, and by Case (b) all even cycles are positive. Thus, (H,Π) has an s-retraction to a single positive edge, contrary to our assumption that the s-core has at least three edges. Hence, (H,Π) must contain a vertex b with a positive loop. (The case where it is negative is handled in the same way.) We claim that (H,Π) also contains a vertex r with a negative loop. Since (H,Π) does not retract to a positive loop, it must contain a negative loop or a negative cycle of length at least 3. Case (b) implies that such a cycle is odd, and Case (c) implies that (H,Π) must contain a negative loop. This establishes the claim.

By the connectivity of H, we can consider a path P from b to r; for simplicity, we switch (H,Π) in such a way that all the edges on P are positive, and let Π' be the obtained signature. Let $P = (r = v_0), v_1, v_2, \ldots, v_{k-1}, (v_k = b)$. Construct $P(H,\Pi')^*$ using the same indicator as in Case (b). Here again, $P(H,\Pi')^*$ has no loop since there is no negative digon. Then, $P(H,\Pi')^*$ contains an odd cycle that goes from $r_0 = v_{0,0}$ to $b_0 = v_{k,0}$ on vertices with even indices, moves to $b_1 = v_{k,1}$, and returns to r_0 on vertices with odd indices, as follows.

$$\begin{aligned} k \text{ even} \quad & v_{0,0}, v_{2,1}, v_{2,0}, v_{4,1}, v_{4,0}, \ldots, v_{k-4,0}, v_{k-2,1}, v_{k,0}, \\ & v_{k,1}, v_{k-1,0}, v_{k-1,1}, \ldots, v_{3,0}, v_{3,1}, v_{1,0}, v_{0,0} \\ k \text{ odd} \quad & v_{0,0}, v_{2,1}, v_{2,0}, v_{4,1}, v_{4,0}, \ldots, v_{k-3,0}, v_{k-1,1}, v_{k,0}, \\ & v_{k,1}, v_{k-2,0}, v_{k-2,1}, \ldots, v_{3,0}, v_{3,1}, v_{1,0}, v_{0,0} \end{aligned}$$

Thus, HOM$(P(H,\Pi)^*)$ is NP-complete. Applying successively Theorem 4.1 and Proposition 2.5, we deduce that EC-HOM$(P(H,\Pi))$ and S-HOM(H,Π) are NP-complete as well. This completes Case (d) and the whole proof. □

Using Proposition 2.5, we can state our dichotomy of Theorem 4.2 in the setting of switching graphs. The (ec-cores of the) switching graphs corresponding to signed graphs on at most two edges are shown in Figure 5.

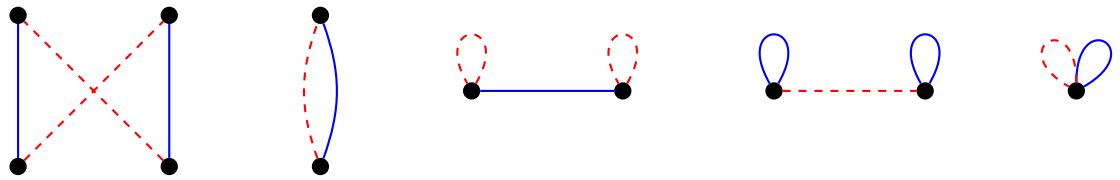

Figure 5: The (ec-cores of) switching graphs constructed from signed graphs on at most two edges.

Corollary 4.3. *Let (H,Π) be a signed graph not containing all three of a negative digon, a positive loop, and a negative loop. If the s-core of (H,Π) has at most two edges, then* EC-HOM$(P(H,\Pi))$ *is polynomial-time solvable; otherwise, it is NP-complete.*

5 Concluding remarks

We conclude the paper with several remarks and questions.

5.1 The smallest case not covered by Theorem 4.2

The cases that are not covered by our dichotomy theorem all contain a negative digon and both kinds of loops. The smallest signed graph that belongs to this family (and is an s-core) is the signed graph (D,Π) on two vertices built from a negative digon by adding to its two vertices a positive loop and a negative loop, respectively (see Figure 6).

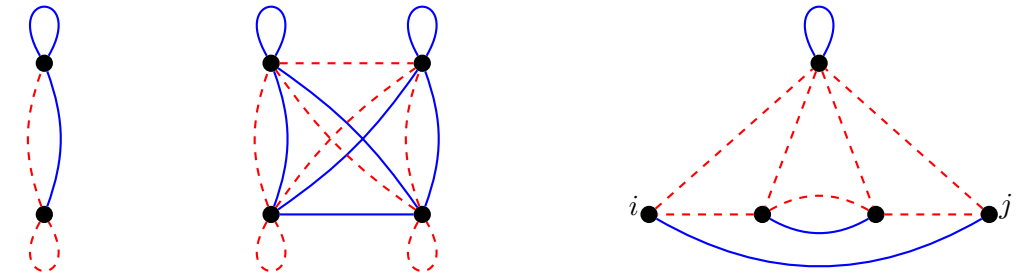

(D,Π) $P(D,\Pi)$ $((I,\Lambda),i,j)$

Figure 6: The signed graph (D,Π), the corresponding switching graph $P(D,\Pi)$, and the indicator $((I,\Lambda),i,j)$.

Proposition 5.1. S-HOM(D,Π) *is NP-complete.*

Proof. The signed graph $P(D,\Pi)$ and the indicator $((I,\Lambda),i,j)$ are shown in Figure 6. The result of the indicator construction $P(D,\Pi)^*$ is the subgraph of $P(D,\Pi)$ induced by the positive edges that are not loops. Hence, it is K_4 minus an edge, that is, the core is K_3. The result follows. □

5.2 List homomorphisms

For a fixed graph H, the *list homomorphism problem* for fixed H, denoted LISTHOM(H), is defined over inputs consisting of pairs (G, L) where G is a graph and L is a list function assigning a list $L(v) \subseteq V(H)$ of possible target vertices to each vertex v of G. The problem asks whether there is a homomorphism $f : G \to H$ such that $f(v) \in L(v)$ for all vertices v of G. The complexity of LISTHOM(H) is known for all undirected graphs H. When H has no loops, Feder, Hell and Huang proved that LISTHOM(H) is NP-complete unless H is bipartite and is the complement of a circular arc graph [6] (then it is polynomial-time solvable). For a signed graph (H, Π), problem s-LISTHOM(H, Π) can be defined analogously as for undirected target graphs, but in the context of s-homomorphisms. We have the following consequence of Theorem 4.2.

Corollary 5.2. *Let (H, Π) be a loop-free signed graph with no negative digon. Then,* s-LISTHOM(H, Π) *is polynomial-time solvable if $\Pi \equiv \emptyset$ (that is, (H, Π) is balanced) and H is a bipartite graph that is the complement of a circular arc graph. Otherwise,* s-LISTHOM(H, Π) *is NP-complete.*

Thus, all polynomial-time cases happen when $\Pi \equiv \emptyset$, and in those cases the dichotomy is described in [6]. When loops are allowed, a dichotomy for the case $\Pi \equiv \emptyset$ is described by Feder, Hell and Huang [7]. We do not know whether in all other cases, s-LISTHOM(H, Π) is again NP-complete.

5.3 Restricted instances

Another line of research is to study s-HOM(H, Π) for special instance restrictions, such as signed graphs whose underlying graph is planar or has bounded degree. Such studies were undertaken for undirected graphs, see for example [10] for bounded degree graphs and [19] for planar graphs.

Acknowledgements

We thank André Raspaud for discussions on the topics of Section 2.4. We also thank Thomas Zaslavsky and the anonymous referees for their detailed comments and helpful suggestions.

References

[1] N. Alon and T. H. Marshall. Homomorphisms of edge-colored graphs and Coxeter groups. *Journal of Algebraic Combinatorics* 8(1):5–13, 1998.

[2] R. C. Brewster. *Vertex colourings of edge-coloured graphs*, PhD thesis, Simon Fraser University, Canada, 1993.

[3] R. C. Brewster. The complexity of colouring symmetric relational systems. *Discrete Applied Mathematics* 49(1–3):95–105, 1994.

[4] R. C. Brewster and T. Graves. Edge-switching homomorphisms of edge-coloured graphs. *Discrete Mathematics* 309(18):5540–5546, 2009.

[5] R. C. Brewster and P. Hell. On homomorphisms to edge-coloured cycles. *Electronic Notes in Discrete Mathematics* 5:46–49, 2009.

[6] T. Feder, P. Hell, and J. Huang. List homomorphisms and circular arc graphs. *Combinatorica* 19(4):487–505, 1999.

[7] T. Feder, P. Hell, and J. Huang. Bi-arc graphs and the complexity of list homomorphisms. *Journal of Graph Theory* 42(1):61–80, 2003.

[8] T. Feder and M. Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: a study through datalog and group theory. *SIAM Journal on Computing* 28(1):57–104, 1998.

[9] F. Foucaud and R. Naserasr. The complexity of homomorphisms of signed graphs and signed constraint satisfaction. Proceedings of the 11th Latin American Symposium on Theoretical Informatics 2014, LATIN 2014, Lecture Notes in Computer Science 8392:526–537, 2014.

[10] A. Galluccio, P. Hell, and J. Nešetřil. The complexity of H-coloring of bounded degree graphs. *Discrete Mathematics* 222(1–3):101–109, 2000.

[11] B. Guenin. Packing odd circuit covers: a conjecture. Unpublished manuscript, 2005.

[12] F. Harary. On the notion of balance of a signed graph. *Michigan Mathematical Journal* 2(2):143–146, 1953-1954.

[13] P. Hell, A. Kostochka, A. Raspaud, and É. Sopena. On nice graphs. *Discrete Mathematics* 234(1–3):39–51, 2001.

[14] P. Hell and J. Nešetřil. *Graphs and Homomorphisms*, Oxford Lecture Series in Mathematics and Its Applications, Oxford University Press, 2004.

[15] P. Hell and J. Nešetřil. On the complexity of H-coloring. *Journal of Combinatorial Theory Series B* 48(1), 92–110, 1990.

[16] T. R. Jensen and B. Toft. *Graph coloring Problems.* Wiley-Interscience Series in Discrete Mathematics and Optimization, 1995.

[17] W. Klostermeyer and G. MacGillivray. Homomorphisms and oriented colorings of equivalence classes of oriented graphs. *Discrete Mathematics* 274(1–3):161–172, 2004.

[18] D. König. *Theorie der endlichen und unendlichen Graphen.* Akademische Verlagsgesellschaft, Leipzig, 1936.

[19] G. MacGillivray and M. Siggers. On the complexity of H-colouring planar graphs. *Discrete Mathematics* 309(18):5729–5738, 2009.

[20] G. MacGillivray and B. Tremblay. Edge-coloured graphs and switching with S_m, A_m, and D_m, in preparation.

[21] G. MacGillivray and J. M. Watten. Colourings of m-edge-coloured graphs and switching, submitted.

[22] A. Montejano, P. Ochem, A. Pinlou, A. Raspaud, and É. Sopena. Homomorphisms of 2-edge-colored graphs. *Discrete Applied Mathematics* 158(12):1365–1379, 2010.

[23] R. Naserasr, E. Rollová, and É. Sopena. Homomorphisms of planar signed graphs to signed projective cubes. *Discrete Mathematics & Theoretical Computer Science* 15(3):1–12, 2013.

[24] R. Naserasr, E. Rollová, and É. Sopena. Homomorphisms of signed graphs. *Journal of Graph Theory* 79(3):178–212, 2015.

[25] P. Ochem, A. Pinlou, and S. Sen. Homomorphisms of 2-edge-colored triangle-free planar graphs. *Journal of Graph Theory*, to appear.

[26] T. Zaslavsky. Characterizations of signed graphs. *Journal of Graph Theory* 5(4):401–406, 1981.

[27] T. Zaslavsky. Signed graphs. *Discrete Applied Mathematics* 4(1):47–74, 1982.

[28] T. Zaslavsky. Signed graph coloring. *Discrete Mathematics* 39(2):215–228, 1982.

[29] T. Zaslavsky. Chromatic invariants of signed graph. *Discrete Mathematics* 42(2–3):287–312, 1982.

[30] T. Zaslavsky. How colorful the signed graph? *Discrete Mathematics* 52(2–3):279–284, 1984.